

PATENT APPLICATION

Methods and Apparatus for Storing and Serving Streaming Media Data

Inventor:

Jason Lango, a citizen of the United States of America, residing at
801 Church St. #1313
Mountain View, CA 94041

Konstantinos Roussos, a citizen of Canada residing at
1063 Morse Ave. Apt. 8-103
Sunnyvale, CA 94089

Robert Tsai, a citizen of the United States of America, residing at
925 Canterbury Rd. #813
Atlanta, GA 30324

J. Christopher Wagner, a citizen of the United States of America residing at
5120 S. Crawford Rd.
Langley, WA 98260-0929

Assignee:

Network Appliance, Incorporated
A California Corporation
495 East Java Drive
Sunnyvale, CA 94089

Entity: Large Entity

Methods and Apparatus for Storing and Serving Streaming Media Data

CROSS-REFERENCE TO RELATED APPLICATIONS

[01] The present invention disclosure claims priority to Provisional U.S. Patent Application Number 60/297,943, filed June 12, 2001, entitled Streaming Media Payload Storage Method and Apparatus. This application is herein by incorporated by reference for all purposes. Co-pending U.S. Patent Application, titled Caching Media Data Using Content-Sensitive Identifiers, filed October 16, 2001, Attorney Docket No.: 020564-000210US, Client Reference No.: P01-1249.2 is also incorporated by reference for all purposes.

BACKGROUND OF THE INVENTION

[02] The present invention relates to data caching. More particularly, the present invention relates to apparatus for caching streaming media and to methods of streaming data from media caches.

[03] Typical file caching methods include a cache receiving a file from a file server, and storing the entire file. Later, when a client desires the file, instead of serving the file from the file server, the file is served from the cache. Because the cache is typically a server that is closer to the client or has higher bandwidth than the file server, the file is served to the client quickly from the cache.

[04] It has been discovered by the inventors, that attempting to apply typical file caching methods to files that include streaming media data, raises many new problems. For instance, serving a streaming media data file from a cache requires much more processing by the cache than with classical file transfers over the web. For example, during normal playback, the cache may need to perform a lot of processing such as packet modification, resequencing, and retiming. As another example, the cache may be called upon to perform random access within the streaming media data file as a result of a client "rewind" or "fast forward" operation. Because, classical caching is typically file-based, such a random access would involve moving within a very large data file.

[05] Another drawback is that since streaming media data files are very large, a huge penalty is incurred if the streaming media data file is deleted. Typically if a file cache determines that it needs more disk space for new files, it will first delete older files, regardless of the size. As an example, if an older file is a streaming media data file that

stores an hour-long program, the entire hour-long program is deleted even if the cache only needs to free up the equivalent of 1 minute of space.

[06] Another drawback is that many different streaming media formats exist, each with its own specific streaming requirements. This is in contrast to classical file transfer over the web, where the files are essentially opaque to the cache and for streaming data to clients, the cache does not need to process the actual contents of the file beyond storage and retrieval.

[07] Thus what is required are improved methods and apparatus for storing and serving streaming media within a cache. Further, what is required are methods and apparatus for providing such solutions in economical ways.

BRIEF SUMMARY OF THE INVENTION

[08] The present invention relates to streaming media caches and methods of operation. More particularly, the present invention relates to efficient storage of streaming media data, and to streaming of media data stored in a streaming media cache to client systems.

[09] In the present disclosure "Streaming media" data generally refers to media intended to be transported at a select (often, subscribed) bit rate, and with a desired timeliness. The streaming media is adapted for playback in a desired order without regard to the order the streaming media data are received by a client system. Streaming media generally conforms to a real-time delivery protocol, such as, e.g., RTSP, RTP, or the like. The media (media clip) represented in the streaming media data may include static images, video data, audio data, executable files, presentation data, applet data, data files, and the like.

[10] The data that is cached in a streaming media cache may be an entire streaming media clip, portions of a streaming media clip, or the like. In the case where there is a streaming media cache hit, the portion of the streaming media stored in the streaming media cache is served to a client. In the case of a streaming media cache miss, the missing portion of a streaming media clip may be retrieved from a media server, instead of an entire streaming media clip. The missing portion of the streaming media clip that is retrieved from an upstream or origin server is then stored and then served to a client. Storing a streaming media clip in the efficient method described below allows the streaming media cache to maintain portions of a streaming media clip that are often requested, and to flush portions of the streaming media clip that are not often requested.

[11] According to one aspect of the invention, a method for storing streaming media data in a cache memory is disclosed. One techniques includes receiving the streaming media data from a streaming media server, the streaming media data comprising a series of packets of

media data, the packets of media data including header data and payload data, separating the header data from payload data, and storing a portion of the header data in a session data object in the cache memory. Other techniques may also include storing the payload data in a first plurality of data objects in the cache memory, wherein each data object of the first plurality of data objects is directly addressable in the cache memory via an associated object handle, and wherein each data object of the first plurality of data objects stores a portion of the payload data.

[12] According to another aspect of the invention, a method for serving streaming media data from a cache memory is disclosed. The method may include retrieving payload data from a first plurality of data objects in the cache memory, wherein each data object of the first plurality of data objects is addressable in the cache memory via an associated object handle, and wherein each data object of the first plurality of data objects stores a portion of the payload data, and retrieving header data from a session data object in the cache memory. Combining the header data and the payload data to form a stream of media data, and serving the stream of media data to a client may also be performed.

[13] According to another aspect of the invention, a computer program product for a computer system including a processor and a memory is described. The memory may include executable code that directs the processor to receive streaming media data from a streaming media server, the streaming media data comprising a series of packets of media data, the packets of media data including header data and payload data, executable code that directs the processor to separate the header data from payload data, and executable code that directs the processor to store a portion of the header data in a session data object in the memory. Additionally executable code that directs the processor to store the payload data in a first plurality of data objects in the memory, wherein each data object of the first plurality of data objects is directly addressable by the processor in the memory via an associated object handle, and wherein each data object of the first plurality of data objects stores a portion of the payload data may also be included in the memory. The memory is typically tangible media such as a hard disk drive, CD-ROM, floppy disk, RAM, and the like.

[14] According to another aspect of the invention, a cache memory is described configured to provide streaming media data. The cache memory may include a session data file storing header data, wherein the header data are selected from the group: encoding scheme, duration, and a plurality of data objects storing payload data, wherein each data object of the first plurality of data objects is addressable in the cache memory via an associated object handle, and wherein each data object of the first plurality of data objects stores a portion of the

payload data. The memory may also include executable code that directs the processor to retrieve header data from the session data object in the cache memory, and executable code that directs the processor to retrieve payload data from the plurality of data objects in the cache memory. Additionally, executable code that directs the processor to combine the header data and the payload data to form a stream of media data, and executable code that directs the processor to serve the stream of media data to a client may be included in the memory.

BRIEF DESCRIPTION OF THE DRAWINGS

[15] Figs. 1A-B illustrate overview diagrams according to embodiments of the present invention;

[16] Fig. 2 is a simplified block diagram of a computer system according to an embodiment of the present invention;

[17] Fig. 3 illustrates a software hierarchy according to embodiments of the present invention;

[18] Figs. 4A-D illustrate a data format hierarchy according to an embodiment of the present invention;

[19] Fig. 5 illustrates a block diagram of a flow chart according to an embodiment of the present invention;

[20] Fig. 6 illustrates a block diagram of a flowchart according to an embodiment of the present invention; and

[21] Fig. 7 illustrates a block diagram of a flowchart according to an embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[22] Fig. 1A illustrates a overview diagram according to an embodiment of the present invention. In particular, Fig. 1A includes a client system 10, a streaming media cache (server) 20, media data server 30 (streaming server), and a router 40. The elements of Fig. 1A are coupled as disclosed over computer networks such as a local area network, wide area networks (Internet), wireless networks or the like.

[23] In one embodiment, client system 10 initially makes a request for a stream of streaming media. The media (media clip) may include static images, video data, audio data, executable files, and the like. This request may take the form of a user clicking upon a URL on a web page, or the like. In this embodiment, this request is intercepted by router 40.

Router 40 may be embodied as a layer 4 or layer 7 switch, a Web Cache Coordination Protocol (WCCP) router, or any other conventional switch or router. In such embodiments, router 40 would be configured to recognize when a request is made by client system 10 for a stream of streaming media.

5 [24] If such a request is determined by router 40, that request is redirected to streaming media cache 20, and not media data server 30. Once streaming media cache 20 receives the request, it makes a determination whether the stream (the entire media clip) or the requested portion of the stream (the request portion of the media clip) has already been cached. If the data has been previously stored, streaming media cache 20 provides the streaming media to
10 client system 10.

[25] In the present embodiment, if the data (requested portion of a stream) has not previously been stored in streaming media cache 20, streaming media cache 20 sends a request to media server 30 for a stream of data including the requested portion of a stream. As the requested portion of the stream of data are delivered to streaming media cache 20, it is forwarded to client system 10, and the portion of the stream of data are stored.

15 [26] For this embodiment, the streaming media traffic is received by media cache 20 from specific ports. In specific embodiments, for RealNetworks RealSystem streaming media, media cache 20 receives streaming media via TCP on port 554; for QuickTime (RTSP) streaming media, media cache 20 receives streaming media via TCP on port 554 and/or via
20 UDP on port 2001; for Microsoft Media Streaming (MMS) streaming media, media cache 20 receives streaming media data via TCP on port 1755; and for HTTP streaming media, media cache 20 receives streaming media data via TCP on port 80, or the like. In other embodiments, other ports for the streaming media may also be used.

[27] The embodiment illustrated above is configured to be accessible from client system
25 10 via a local area network. It should be understood that streaming media cache 20 may be alternatively positioned at other points in the network, for example, at the edge of a point of presence network on the Internet, and the like. An example is illustrated in Fig. 1B

[28] Fig. 1B illustrates a overview diagram according to another embodiment of the
30 present invention. In particular, Fig. 1B includes a client system 15, a streaming media cache (server) 25, media data server 35 (streaming server), and a router 42. The elements of Fig. 1B are coupled as disclosed over computer networks such as a local area network, wide area networks (Internet), wireless networks or the like. In this embodiment, streaming media cache 25 may be embodied as an accelerator on the edge of a point of presence (POP).

[29] In this embodiment, client system 15 initially makes a request for a stream of streaming media (representing a streaming media clip). This request may take the form of a user clicking upon a URL on a web page, or the like. In this embodiment, the request is passed over the wide area network and is intercepted by router 42. Router 42 may be embodied as a layer 4 or layer 7 switch, a WCCP router, or any other conventional switch or router. In this embodiment, router 42 would be configured to recognize when a request is made by client system 10 for a stream of streaming media.

[30] If such a request is determined by router 42, that request is redirected to streaming media cache 25, and not media data server 35. Once streaming media cache 25 receives the request, it makes a determination whether the streaming media clip or the requested portion of the streaming media clip has already been cached. If the data has been previously stored, streaming media cache 25 provides the streaming media to client system 10.

[31] In the present embodiment, if the data has is not stored in streaming media cache 25, streaming media cache 25 sends a request to media server 35 for the missing data. As the stream of data (including the portion of the streaming media clip) is delivered to streaming media cache 25, it is forwarded to client system 15. The missing portion of the streaming media clip is then stored in streaming media cache 25. Details of the storage format and the process of storing and retrieving the stream of data are described in greater detail below.

[32] For this embodiment, the streaming media traffic is sent by media cache 20 to specific ports. In specific embodiments, for RealSystem streaming media, media cache 20 sends streaming media via TCP on port 554; for QuickTime (RTSP) streaming media, media cache 20 sends streaming media via TCP on port 554 and/or via UDP on port 2001; for Microsoft Media Streaming (MMS) streaming media, media cache 20 sends streaming media data via TCP on port 1755; and for HTTP streaming media, media cache 20 sends streaming media data via TCP on port 80, or the like. In other embodiments, other ports for the streaming media may also be used.

[33] In other embodiments of the present invention, one or more streaming media caches may be positioned simultaneously at the illustrated locations between client system 15 and media server 35. Additional streaming media caches may also be positioned at other locations between client system 15 and media server 35, for example at a user ISP, on an intranet, and the like. In light of this disclosure, it will be apparent that many other network configurations can incorporate embodiments of the present invention.

[34] Fig. 2 is a simplified block diagram of a computer system 45 according to an embodiment of the present invention. Computer system 45 may be used as client system 10,

streaming media cache 20, and/or media data server system 30. Computer system 45 may be a stand-alone computer system, a computer “appliance,” or the like.

[35] As shown in Fig. 2, computer system 45 includes at least one processor 50, which communicates with a number of peripheral devices via a bus subsystem 55. These peripheral devices may include a storage subsystem 60, comprising a memory subsystem 65 and a file storage subsystem 70 user interface input devices 75, user interface output devices 80, and a network interface subsystem 85. The input and output devices allow user interaction with computer system 45. A user may be a human user, a device, a process, another computer, and the like.

[36] Network interface subsystem 85 provides an interface to other computer systems. Embodiments of network interface subsystem 85 include an Ethernet card, a modem (telephone, satellite, cable, ISDN), (asynchronous) digital subscriber line (DSL) units, and the like. Network interface 250 is coupled to a typical network as shown.

[37] User interface input devices 75 may include a keyboard, pointing devices such as a mouse, trackball, touchpad, or graphics tablet, a scanner, a barcode scanner, a touchscreen incorporated into the display, audio input devices such as voice recognition systems, microphones, and other types of input devices. In general, use of the term “input device” is intended to include all possible types of devices and ways to input information using computer system 50.

[38] User interface output devices 80 may include a display subsystem, a printer, a fax machine, or non-visual displays such as audio output devices. The display subsystem may be a cathode ray tube (CRT), a flat-panel device such as a liquid crystal display (LCD), or a projection device. The display subsystem may also provide non-visual display such as via audio output devices. In general, use of the term “output device” is intended to include all possible types of devices and ways to output information from computer system 45.

[39] Storage subsystem 60 may be configured to store the basic programming and data constructs that provide the functionality of the computer system and of the present invention. For example, according to an embodiment of the present invention, software modules implementing the functionality of the present invention may be stored in storage subsystem 60. These software modules may be executed by processor(s) 50 of computer system 45. In a distributed environment, the software modules may be stored on a plurality of computer systems and executed by processors of the plurality of computer systems. Storage subsystem 60 may also provide a repository for storing various databases that may be used to store information according to the teachings of the present invention. For example, a cache entry

hash table, discussed below, may be stored in storage subsystem 60 of media server 30.

Storage subsystem may also function as a cache of streaming media cache 20. Storage subsystem 60 may comprise memory subsystem 65 and file storage subsystem 70.

[40] Memory subsystem 65 may include a number of memories including a main random access memory (RAM) 90 for storage of instructions and data during program execution and a read only memory (ROM) 95 in which fixed instructions are stored. RAM 90 is typically also used for execution of programs, storage of data, and the like.

[41] File storage subsystem 70 provides persistent (non-volatile) storage for program and data files, and may include a hard disk drive, a floppy disk drive along with associated removable media, a Compact Digital Read Only Memory (CD-ROM) drive, an optical drive, removable media cartridges, and other like storage media. One or more of the drives may be located at remote locations on other connected computers.

[42] A memory buffer 97 is also provided in storage subsystem 60. In this embodiment, memory buffer 97 is a special buffer memory coupled to file storage subsystem 70. More specifically, memory buffer 97 provides a temporary storage area for data retrieved from and data sent to file storage subsystem 70. Memory buffer 97 may also provide a temporary storage area for data received from a streaming media server (or other upstream server) and for data to be sent to client systems. As will be discussed below, the type of data may include streaming media payload data.

[43] In the present embodiment, computer system 45 typically also includes software that enables it to send and receive data and communications to and from client systems 10 and media data server 30 using communications protocols including, HTTP, S-HTTP, TCP/IP, UDP, SSL, RTP/RTSP and the like. In alternative embodiments of the present invention, other software and transfer and communication protocols may also be used, for example IPX, UDP or the like.

[44] Bus subsystem 55 provides a mechanism for letting the various components and subsystems of computer system 45 communicate with each other as intended. The various subsystems and components of computer system 45 need not be at the same physical location but may be distributed at various locations within a network. Although bus subsystem 55 is shown schematically as a single bus, alternative embodiments of the bus subsystem may utilize multiple busses.

[45] Computer system 45 itself can be of varying types including a personal computer, a portable computer, a workstation, a computer terminal, a network computer, a mainframe, a kiosk, a personal digital assistant (PDA), a wireless communication device such as a cell

phone, an entertainment console (PS2, X-box) or any other data processing system. Due to the ever-changing nature of computers and networks, the description of computer system 45 depicted in Fig. 1B is intended only as a specific example for purposes of illustrating an embodiment of the computer system.

5 [46] In one embodiment, computer system 45 is embodied as a network cache (appliance) in a product called "NetCache" available from NetworkAppliance, Incorporated. The NetCache family of products currently includes the NetCache C1100, NetCache C3100, and NetCache C6100 including proprietary, but available hardware and software. Embodiments of the present invention may also be implemented in future additions to the NetCache family of products.

10 [47] It will be readily apparent to one of ordinary skill in the art that many other hardware and software configurations are suitable for use with the present invention. For example, other types of processors are contemplated, such as the Athlon™ class microprocessors from AMD, the Pentium™ -class or Celeron™-class microprocessors from Intel Corporation, PowerPC™ G3 or G4 microprocessors from Motorola, Inc., Crusoe™ processors from Transmeta, Inc. and the like. Further, other types of operating systems are contemplated in alternative embodiments including WindowsNT™ from Microsoft, Solaris from Sun Microsystems, LINUX, UNIX, MAC OS X from Apple Computer Corporation, BeOS™, and the like. Many other configurations of a computer system are possible having more or fewer components than the computer system depicted in Fig. 1B.

15 [48] Fig. 3 illustrates a software hierarchy according to embodiments of the present invention. In particular, Fig. 3 includes a three-tiered hierarchy including an operating system level (layer) 100, a data handling level (layer) 110, and a protocol level (layer) 120.

20 [49] In the present embodiment, as illustrated, operating system level(layer) 100 includes portions of the Berkeley Software Distribution (BSD) operating system. Additionally, operating system level 100 includes software provided by the assignee of the present invention: Data ONTAP™, a Network Appliance brand operating system with Write Anywhere File Layout (WAFL™), a Network Appliance brand file system. In the present embodiment, the Data ONTAP™ operating system provides efficient file service by using
25 file-system technology and a microkernel design geared towards network data access. The WAFL™ file system provides efficient file storage and retrieval based upon efficient access algorithms and data structures. Additionally, network communications using Transmission Control Protocol (TCP) and UDP are also supported at operating system level 100. Of course
30 other types of operating systems can also be used.

[50] As illustrated in Fig. 3, data handling level(layer) 110 includes a packet pacing subsystem (SMPACER) 130 and a streaming disk subsystem (SMDISK) 140. In the present embodiment, streaming disk subsystem 140 is used to retrieve data packets from the file system and to provide the data to SMPACER 130. As will be described below, in one embodiment, SMDISK 140 receives streaming media data packets and in turn SMDISK 140 creates a series of specialized data objects for storing the data. Further, SMDISK 140 receives the specialized data objects from the file system and stores the data packets into a buffer for output as streaming media.

[51] In this embodiment, SMPACER 130 receives data packets (meta-data and payload data) via a pointer to a buffer location or the like from SMDISK 140. In turn, SMPACER 130 sends the pointers to protocol level(layer) 120. As described below, protocol level 120 formats the packets according to the desired streaming protocol. The formatted streaming packets are then received by SMPACER 130. Based upon delivery times for each packet, SMPACER 130 then sends a stream of packets to the client system at the desired rate. In particular, protocol level 120 “filters” or adjusts the “delivery time” of packets to be output to clients, and the like. The adjusted meta-data and the payload data are then output by SMPACER 130 to a client, based upon the adjusted delivery time.

[52] In this embodiment, protocol level 120 includes support for at least one, but typically for more than one streaming media protocols. The support includes encoding of data to form streams of streaming media and decoding of streams of streaming media. In one example, a streaming media protocol is the Microsoft Media Streaming (MMS) protocol. By supporting the MMS protocol, streams of MMS formatted data can be received from a streaming media (upstream or origin) server and the streamed (payload) data can be retrieved. This payload data can be sent to data handling layer 110 via SMDISK 140 for storage. Additionally, payload data determined by SMDISK 140 can be encoded into streams of MMS data. The encoded data are then sent to SMPACER 130 for paced delivery to a client system. The client system may play the encoded data via a player such as Microsoft Windows Media Player, and the like.

[53] In another example, a streaming media protocol is the Real Time Streaming Protocol (RTSP). In addition to RTSP support, one embodiment includes Apple QuickTime format support and RealNetworks RealSystem format support. By supporting these protocols, streams of QuickTime formatted data or RealSystem data can be received from streaming media servers and the respective streaming (payload) data are retrieved. These payloads are then sent to data handling layer 110 via SMDISK 140 for storage. Additionally, payload data

from SMDISK 140 can be encoded into streams of data and delivered to the client by SMPACER 130. The streaming data can be played on client systems via a QuickTime player or a RealSystem player, and the like. In other embodiments, other types of streaming media encoding schemes may be supported.

5 [54] The above hierarchy has been described in embodiments as being implemented via software. However, it should be understood that some functions may be implemented in hardware or firmware. Accordingly, additional embodiments of the above may be implemented via hardware, firmware, software, and combinations thereof. Further description of SMPACER 130 will be given below.

10 [55] Figs. 4A-D illustrate a data format hierarchy according to an embodiment of the present invention. In particular, Figs. 4A-D illustrate an internal storage structure / format used by embodiments for storing data that will be streamed to client systems.

15 [56] An example of a streaming media cache implementing a data storage structure described below is a NetCache™ streaming media cache. NetCache™ (latest version 5.2) includes a combination of hardware and software available from the assignee of the present patent application. Embodiments of the present invention may stream data to client systems in a variety of streaming media protocols, including Microsoft Media Streaming (MMS) protocol used by Windows Media Player™; Real Time Streaming Protocol (RTSP) used by Quicktime™ from Apple Corporation and RealSystem™ from RealNetworks; and the like.

20 [57] As illustrated in Fig. 4A, the present embodiment includes a cache entry table hash table 200 and a plurality of entries wherein each entry includes an object identifier 210. In one embodiment, object identifiers are file names that have been hashed. Further details regarding this aspect of the invention are disclosed in the co-pending application cited above. Cache entry table 200 typically also includes a plurality of object handles 220 for a particular object. In the present embodiment, object handles 220 may be a reference or pointer to an object 230 corresponding to the object identifier and stored in a cache 235.

25 [58] In the present embodiment, object handles 220 may be used to retrieve the corresponding object 230 from cache 235. According to an embodiment of the present invention, objects 230 are stored as separate data files in cache 235. In this embodiment, each object handle 220 corresponds to a file handle and the object itself is stored as a file. Accordingly, the individual files are each independently accessible in cache 235 by a file system.

30 [59] Fig. 4B illustrates a session description 250 (stored in a session data file or session data object) and logical streams of data 260 and 270 according to an embodiment. Logical

stream 260 represents data for streaming media encoded in a first encoding scheme and logical stream 270 represents data for streaming media encoded in a second encoding scheme.

[60] In the present embodiment, each of the encodings of the data are considered separate streams of data and are stored separately. This is in contrast to cases where multiple encodings of a data stream are packaged and stored within a single data file. An example of the latter is used by RealNetworks. In particular, a data file used by RealSystem may include an encoding of data destined for 56Kbps clients, and an encoding of data destined for 384 Kbps clients. In the present embodiment, the encoding of data destined for different bit rate clients would be stored separately. For example, a 56 Kbps encoding would be stored in logical stream 260 and a 384 Kbps encoding would be stored in logical stream 270. Other typical types of parameters that may be varied for different encodings may include the bit rate, the content (e.g. abridged, unabridged), the media type (audio and/or video), thinning parameters (frame dropping), and the like.

[61] In Fig. 4B, session description (stored in a session data object or session data file) 250 may include a description of the various streams of data stored in logical streams 260 and 270. The description may include an enumeration of the various encoding schemes (e.g. 56 Kbps, 128 Kbps, ISDN), copyright and authoring data, presentation or play-time(duration) of the stream, version data, and the like.

[62] As an example, a sample session description for RTSP is as follows. In particular, it illustrates extrinsic properties of the media file (author, title, copyright), as well as intrinsic properties of the media file (number of media tracks, length of media file, encoding bitrate, MIME type, and codec of each media track, etc.). All of this data together serves to help uniquely identify a particular version of the URL used to access the streaming media file.

```
v=0
o=- 983139433 983139433 IN IP4 172.30.200.154
s=G2 Video Experience
i=RealNetworks ©1998
t=0 0
a=SdpplinVersion:1610642970
a=Flags:integer;2
a=IsRealDataType:integer;1
a=StreamCount:integer;2
a=Title:buffer;"RzIgVmlkZW8gRXhwZXJpZW5jZQA="
```

```

a=Copyright:buffer;"qTE5OTgA"
a=Author:buffer;"UmVhbE5ldHdvcmtzAA=="
...
a=range:npt=0-0
5 m=audio 0 RTP/AVP 101
b=AS:6
a=control:streamid=0
a=range:npt=0-59.773000
a=length:npt=59.773000
10 a=rtptime:101 x-pn-realaudio
a=mimetype:string;"audio/x-pn-realaudio"
a=MinimumSwitchOverlap:integer;200
a=StartTime:integer;0
a=AvgBitRate:integer;6000
15 a=EndOneRuleEndAll:integer;1
a=AvgPacketSize:integer;288
a=SeekGreaterOnSwitch:integer;0
a=Preroll:integer;4608
a=MaxPacketSize:integer;288
20 a=MaxBitRate:integer;6000
a=RMFF 1.0 Flags:buffer;"AAQAAgAAAAIAAA=="
...
a=StreamName:string;"audio/x-pn-multirate-realaudio logical
stream"
25 ...
m=video 0 RTP/AVP 101
b=AS:50
a=control:streamid=1
a=range:npt=0-57.333000
30 a=length:npt=57.333000
a=rtptime:101 x-pn-realvideo
a=mimetype:string;"video/x-pn-realvideo"
a=MinimumSwitchOverlap:integer;0
a=StartTime:integer;0

```

```

a=AvgBitRate:integer;50000
a=EndOneRuleEndAll:integer;1
a=AvgPacketSize:integer;538
a=SeekGreaterOnSwitch:integer;1
a=Preroll:integer;5707
a=MaxPacketSize:integer;607
a=MaxBitRate:integer;50000
a=RMFF 1.0
Flags:buffer;"AAoAAgAAAAAAgACAAAAAgAAAAIAAA=="
...
a=StreamName:string;"video/x-pn-multirate-realvideo logical
stream"
...

```

[63] In the present embodiment, logical streams of data, such as logical stream 260 is made up of a series of data objects 280. As described in Fig. 4A, data objects 280 are physically separate files that are directly accessible, independent of other data objects 280, through use of cache entry hash table 200. In this embodiment, data objects 280 together store the “media payload” provided by streaming media encoded in a given encoding scheme. For example, the media payload may include the actual media data included in streaming media packets for a 56 Kbps stream, or the like. More particularly, data objects 280 store the media payload that has been converted from the format in which the origin server stores the media data into the network format for transmission to the client system and the cache. Accordingly, the data objects include data that are optimized for delivery to the client system (e.g., encapsulated in network protocol).

[64] In the present embodiment, each data object 280 is used to store data having an associated and/or a predetermined amount of play time (duration). That is, each data object 280 is used to store media payload data that will be output as streaming data that will be played on a client system for a specific amount of time or duration. For example, in one embodiment, each data object 280 is used to store data that will be streamed to a client as 20 seconds of a music stream, video stream, or the like. In other embodiments, each data object 280 may store a media payload (data) having different duration, such as less than or equal to approximately 5 seconds, 10 seconds, 20 seconds, 30 seconds, 1 minute, or the like.

[65] In one embodiment of the present invention, the duration of output for the media payload stored in typical data objects may be fixed for each data object among logical streams 260 and 270 (e.g. 15 seconds of a stream). However, in other embodiments, the duration of output for the media payload stored in typical data objects in logical stream 260 and data objects in logical 270 may be different. For example, for logical stream 260, the duration may be 15 seconds per data object, and for logical stream 270, the duration may be 30 seconds per data object, and the like.

[66] In another embodiment, each data object 280 may store specific amounts of data instead of a specific duration for data. For example, each data object 280 may store a predetermined number of bytes of data, for example, less than or equal to approximately 64 Kbytes, 128 Kbytes, 512 Kbytes, 1 Mbyte, or the like. In another embodiment, each data object 280 may simply store “chapters” or logical segments of a movie or video, and the like. In one embodiment, each data object 280 stores a fixed number of data chunks, as described below.

[67] In one embodiment of the present invention, data objects 280 store non-overlapping data, or unique portions of the media data. That is, each of the data objects 280 may be configured to store a portion of the media data that is unique to a reference (e.g., URL) in the request to locations in the origin (or upstream) server at which the media file is stored. In another embodiment, data objects 280 may store media data that overlaps or is redundant.

[68] Fig. 4C illustrates a detailed description of a data object according to an embodiment of the present invention. As illustrated, Fig. 4C, a data object 300 includes object meta-data portion 310, and data chunks 320.

[69] In this embodiment, object meta-data portion 310 is used to store data about data object 300. Such meta-data, or header data, may include file format version numbers, the number of data chunks 320 stored, the beginning presentation time and ending presentation time for data objects, and the like. In other embodiments, additional data may be stored in object meta-data portion 310 such as the data object number, protocol-specific per-data object data, a total number of bytes of payload and meta-data per data object, the number of data packets per data object, any end of stream indicators, checksum bits and the like.

[70] In one embodiment, each data chunk 320 is also used to store data of a predetermined amount of presentation or play time (duration). That is, each data chunk 320 is used to store streaming data that will be played on a client system for a specific amount of time. For example, in one embodiment, each data chunk 320 is used to store 20 seconds of a music stream. In other embodiments, each data chunk 320 may store data having different duration,

such as less than or equal to approximately 5 seconds, 10 seconds, 20 seconds, 30 seconds, 1 minute, or the like. In one embodiment of the present invention, the duration may be fixed for each data chunk 320 within data object 300. However, in other embodiments, data objects may have different durations.

5 [71] In another embodiment, each data chunk 320 may store specific amounts of data. For example, each data chunk 320 may store a predetermined number of bytes of data, for example, less than or equal to approximately 32 Kbytes, 64 Kbytes, 128 Kbytes, 512 Kbytes, 1 Mbyte, or the like. In one embodiment, each data chunk has a fixed number of data packets. In still other embodiments, data chunks 320 may have a varying number of data
10 packets.

[72] As will be described below, in the present embodiment, each data chunk 320 is used to store the actual streaming media data. More particularly, each data chunk 320 is used to store packets of data that will be streamed to a client system.

[73] Fig. 4D illustrates a detailed description of a data chunk according to an embodiment of the present invention. Each data chunk 340 includes a chunk meta-data portion 350,
15 packet meta-data 360, packet match bits 370, and packet payloads 380.

[74] In this embodiment, chunk meta-data portion 350 is used to store data about data chunk 340. For example, chunk meta-data portion 350 may specify the number of packet payloads (packets) 380, a file offset for a previous data chunk within the same data object, a file offset for the next data chunk within the same data object, the number of data packets in a data chunk, compressed packet meta-data for the packets, described below, and the like. In
20 additional embodiments, the data chunk meta-data header may also include packet meta-data for all the data packets including the duration (playback duration) of the payload, the presentation time of the payload (e.g. time within a movie), the delivery time of the payload (a time SMPACER 130 delivers the payload data to the client), protocol-specific data of the payload, and the like. Other types of data may be stored in chunk meta-data portion 350 in
25 other embodiments, such as timing information, and the like.

[75] Payload packets 380 are used to store streaming data packets that make up the streaming media. For example, payload packets 380 may store audio data, image data,
30 audiovisual data, and the like. As will be described below, the streaming data packets may be received as stream of data from a streaming media server, or may be derived from a data file received from the streaming media server. For Windows Media Player streaming media, payload packets 380 range from 200 bytes to 18 Kbytes of data, and for RealSystem streaming media and QuickTime streaming media, packet payloads 380 range from

approximately 200 to 1.5 Kbytes, typically 600 bytes. The number of packet payloads in data chunk 340 typically depends upon the size of packet payloads 380.

[76] In this embodiment, packet meta-data 360 is used to store information relevant to or associated with each payload packet 380. Types of information may include the delivery time and the presentation time, file offset of the respective payload packet 380, and the like. In the present example, the delivery time is the time SMPACER 130 should send the packet payload to the client. In contrast, the packet presentation time is the time within the media stream that the payload is displayed by the client system.

[77] Packet match bits 370 are used in the present embodiment to store information specific to the streaming media protocol. For example, packet match bits 370 may store data such as flags to identify the start of video key-frames, such as I, B, and or P key frames, or the like. In this embodiment, packet match bits 370 are used to determine the first sendable payload (keyframe) that satisfies a SEEK request by the client system. In one embodiment, the match bits may be embodied as single bit, however, in other embodiments of the present invention, additional match bits may be used to represent any number of criteria, for example, for selecting which packet will be delivered first, and the like. For Windows Media Player streaming media, packet match bits 370 may be as small as a single bit, and for RealSystem streaming media and QuickTime streaming media, packet match bits 370 are approximately 32 bits.

[78] In this embodiment, the match bits are logically grouped together and separated from the remaining packet metadata. By grouping of the match bits together, the match bits can be compressed into, for example, a single word, thereby saving memory space.

[79] Such key frame data are useful when a client system requests to move around the stream data, for example, when jumping to a particular presentation time T within the stream.

In this embodiment, based upon packet match bits 370, the key frame immediately before presentation time T is retrieved and the play-back is begun from that key frame. It has been discovered that in one embodiment, playing-back stream data from the immediately-preceding keyframe reduces the amount of media artifacts or blank time of the stream when played on the client system.

[80] Fig. 5 illustrates a block diagram of a flow chart according to an embodiment of the present invention. More particularly, Fig. 5 illustrates a process of storing streaming media data into embodiments of the present invention. In the below embodiments, data are typically stored using the data format hierarchy illustrated in Figs. 4A-D.

[81] In Fig. 5, a data packet delivered from a streaming media server is received in step 400. In one embodiment, the streaming media server streams data to the streaming media cache. In such a case, a packet of data from the stream includes header data and a data packet (payload data). In another embodiment, the streaming media server sends a data file including the entire data stream to the streaming media cache. In this case, data packets and the header data are buried within the data file.

[82] In the present embodiment, packet meta-data for a particular packet of data are then identified, step 410. In one embodiment of the present invention, the packet meta-data are derived from the header data of a particular data packet. In another embodiment, the packet is derived from the data file. The packet meta-data may include a presentation time for a data packet, an indication of a video key-frame, and the like. In this example, presentation time is the time within a media stream where the data packet is presented, for example, a data packet may have a presentation time of 20.5 seconds to 20.6 seconds representing when the data packet is output on the client system.

[83] Next, a determination is made as to whether a new data object should be created, step 420. A new data object is typically created when a first data packet is received, or as described below a previous data object is full. In one embodiment, a new data object is created, step 430.

[84] Next, a determination is made as to whether a new data chunk within the data object should be created, step 440. A new data chunk is typically created when a first data packet is received, or as described below, a data chunk is closed after including the previous data packet. In one case a new data chunk is created, step 450.

[85] The data packet and the packet meta-data are then typically written to a buffer location in the streaming media cache random access memory, step 460. This buffer may be RAM 90 or buffer 97. In this embodiment, it is then determined whether the data packet is the last one for a given data chunk, step 470. If not, the process above is repeated for the next data packet.

[86] When the data chunk is full, the chunk meta-data are determined, and the data chunk is written to random access memory (or to disk memory), step 480. In this embodiment, it is then determined whether the data chunk is the last one for a given data object, step 490. If not, the process above is repeated for the next data packet.

[87] In this embodiment, when the data object is full, the object meta-data described above is determined, and the data object is written to disk memory, step 400. The process above may then be repeated until there are no more data packets in the media stream.

[88] Accordingly, using the above steps, streaming media data may be received by a streaming media cache and stored in a disk memory in the object-based scheme described above. Additionally, streaming media data may be received in the form of a data file. This data file is parsed and the data are also stored in a disk memory in the object-based scheme described above.

[89] In the above embodiment, most of the functions are performed by SMDISK 140, discussed in Fig. 3. In particular, steps 400 and 470 are typically performed at least in part by SMDISK 140; and step 480 is typically performed by a file system within operating system level 100.

[90] Fig. 6 illustrates a block diagram of a flowchart according to an embodiment of the present invention. In particular, Fig. 6 illustrates an overview of the process of retrieving data stored in a disk memory of the streaming media cache as described in Figs. 4A-D and forming a stream of streaming media data to a client system.

[91] In this example, a client system requests streaming media from an embodiment of a streaming media cache, step 500. In one case, a request for streaming media may be made directly from a client system or via a proxy. Such a request is typically in the form of a URL, or the like. Additionally, the request may specify a presentation time T that represents the time where the playback should begin. Most requests set T equal to zero, however T is typically non-zero when the client system jumps around the media stream (e.g. makes a "seek" request).

[92] If the client system does not terminate its connection with the streaming media cache, step 510, a determination is made as to whether to playback the streaming data or not, step 520. In embodiments of the present invention, other types of client events may be specified, such as disconnecting, a play request, a pause request, a stop request, a seek request, notification to the cache that while the client is receiving streaming data from the cache, that a future object is missing and needs to be prefetched, and the like.

[93] In the present embodiment, if streaming data are to be streamed to the client system, the presentation time T is determined, step 530. Next, based upon the time T, the payload packet that includes data having the presentation time T is located, step 540. This step is typically performed in part by SMDISK 140. Next, the data are then formatted for the specific protocol and then sent to the client system, step 550. This step is typically performed in part by SMPACER 130 and protocol level 120. More detailed descriptions of the above steps is given below.

[94] Fig. 7 illustrates a block diagram of a flowchart according to an embodiment of the present invention. In particular, Fig. 7 illustrates a more detailed process of locating and serving data.

[95] In the present embodiment, in response to the presentation time T, the streaming media cache initially determines which data object to retrieve first, step 600. In the embodiment above, because an amount of time for each data object is fixed, for example at 10 seconds, the appropriate data object can easily be determined. For example, if the presentation time T were 5 minutes into a data stream, the appropriate data object would be the thirtieth one $((5 \text{ minutes} \times 60 \text{ seconds} / \text{minute}) / 10 \text{ seconds} / \text{data object} = 30)$. In one embodiment, the URL of the file, along with the presentation time T is first hashed, and the hash is then used to access the cache entry hash table illustrated in Fig. 4A. In another embodiment, a URL of the file, the type of encoding of the file, a validator for the file, and the time T is hashed, and the hash is used to access the cache entry hash table illustrated in Fig. 4A. In return, the cache entry hash table provides the appropriate file handle of the targeted data object.

[96] Based upon the file handle, the object meta-data are first retrieved, step 610. The data are typically stored within RAM 90. Based upon the number of chunks of data within the target data object, the target data chunk is determined. In the present embodiment, the meta-data of the first data chunk in a data object is first retrieved and stored within RAM 90. This data also includes the packet meta-data for that data chunk. Then, using the chunk meta-data, by using the file offset meta-data, the target data chunk containing the desired packet payload (keyed by presentation time) is determined.

[97] Next, the chunk meta-data of the target data chunk is retrieved, step 620. The chunk meta-data are stored within RAM 90 for access by processor 50. As described above, the chunk meta-data may specify the number of payload packets stored within the chunk. Next, based upon the number of payload packets within the data chunk, the target payload packet is determined. The packet meta-data of the target payload packet is then retrieved and stored within RAM 90 for access by processor 50, step 630.

[98] In the present embodiment, packet match bits 270 are also retrieved, and if compressed, uncompressed. The packet match bits 270 are typically stored within RAM 90.

[99] In the present embodiment, portions of the packet meta-data and the target payload packet are then combined, step 640. The resulting packet is sent to the client system, step 650. In embodiments of the present invention, the target payload packet is the same as what was received from the origin server. Further, the packet meta-data are typically protocol-

specific header data, i.e. the data depends upon the type of stream provided, such as Quicktime, Windows Media, and the like. for example, the meta-data may include a per-client sequence number, packet timing information, and the like.

[100] After this target payload packet is sent, this embodiment attempts to iterate to the next payload packet, step 660. If the target payload packet is the last one of the target data chunk, step 670, this embodiment attempts to iterate to the next data chunk. If the target data chunk is the last one of the target data object, step 680, this embodiment attempts to iterate to the next data object. If the target data object is the last one of the stream, step 690, the stream terminates.

[101] In the above embodiment steps 600-630 are performed at least in part by SMDISK 140; step 640 is performed at least in part by SMPACER 130; and step 650 is performed at least in part by SMPACER 130. More specifically, SMDISK 140 typically retrieves packet meta-data and packet payloads from the cache memory (hard disk) and stores them into a memory buffer, such as buffer 97. SMDISK 140 then gives pointers to these buffer locations to SMPACER 130, and in turn SMPACER 130 gives the pointers to these buffer locations to protocol level 120. An encoding protocol in protocol level 120 processes the meta-data portion, and importantly, then simply appends the packet payload to form an encoded packet. This encoded packet is sent to SMPACER 130 for paced delivery to a client.

[102] As illustrated above, packet payloads are simply stored and retrieved from the cache memory (hard disk) and no processing occurs on such data. The payload data are merely segmented into convenient-sized data chunks and data objects by SMDISK 140 and then stored in the cache memory. As discussed above, these data objects are individually accessible on a file level.

[103] Advantages to embodiments of the present invention are believed to include that partial media streams can be saved within a cache. As disclosed above, as objects are populated, they are written to disk. Thus for example, when a first client system receives a data stream from a media server via the streaming media cache, the streaming media cache stores the data to disk. If the first client system terminates the stream half way through the stream, the streaming media cache stores half the stream. Later, the streaming media cache can serve the first half of the stream from cache.

[104] Another advantage is that when the streaming media cache is full, it does not have to eject or flush an entire data stream, but can simply eject less frequently requested data objects. For example, the streaming media cache can eject the last 10 minutes of an hour long stream, while maintaining storage of the first 50 minutes. Thus, in contrast to

conventional file level caching, embodiments of the present invention have an advantageous level of control over streaming media data. When additional space is required in the cache, seldom accessed portions of a data stream may thus be ejected freeing up space in the media cache for newer media data, or the like. In other embodiments, other indicia for determining when to flush portions of data streams may include upon number of requests for data, last request time for the data, based upon priority for data, and the like.

[105] In view of the above disclosure, many other variations can be envisioned. For example, the data hierarchy can simply be modified according to engineering requirements. As an example, the granularity the data hierarchy can easily be changed and varied based upon experimental data, and the like. As another example, the sizes of the data structures may also be varied according to engineering requirements.

[106] The invention has been described in embodiments above as a file cache or a streaming media cache. It should be understood, however, that, embodiments may be embodied in any computer system as a stand-alone system, or as part of another system. For example, one embodiment may be integrated into a computer system that includes web server software, database software, and the like. As another example, one embodiment may be distributed among a set of computer systems in a network, or the like. In similar examples, when there is a miss, embodiments of the present invention may access other embodiments in the network before attempting to access an origin server, or the like.

[107] In other embodiments of the present invention, combinations or sub-combinations of the above-disclosed invention can be advantageously made. The block diagrams of the architecture and flowcharts are grouped for ease of understanding. However it should be understood that combinations of blocks, additions of new blocks, re-arrangement of blocks, and the like are contemplated in alternative embodiments of the present invention.

[108] The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.